

Initial Characterization of Optical Communications With Disruption-Tolerant Network Protocols

Joshua Schoolcraft

Communications Networks Group
JPL, California Institute of Technology
Pasadena, CA, USA

Keith Wilson

Optical Communications Group
JPL, California Institute of Technology
Pasadena, CA USA

Abstract—Disruption-tolerant networks (DTNs) are groups of network assets connected with a suite of communication protocol technologies designed to mitigate the effects of link delay and disruption. Application of DTN protocols to diverse groups of network resources in multiple sub-networks results in an overlay network-of-networks with autonomous data routing capability. In space environments where delay or disruption is expected, performance of this type of architecture (such as an interplanetary internet) can increase with the inclusion of new communications mediums and techniques. Space-based optical communication links are therefore an excellent building block of space DTN architectures. When compared to traditional radio frequency (RF) communications, optical systems can provide extremely power-efficient and high bandwidth links bridging sub-networks. Because optical links are more susceptible to link disruption and experience the same light-speed delays as RF, optical-enabled DTN architectures can lessen potential drawbacks and maintain the benefits of autonomous optical communications over deep space distances. These environment-driven expectations – link delay and interruption, along with asymmetric data rates – are the purpose of the proof-of-concept experiment outlined herein. In recognizing the potential of these two technologies, we report an initial experiment and characterization of the performance of a DTN-enabled space optical link. The experiment design employs a point-to-point free-space optical link configured to have asymmetric bandwidth. This link connects two networked systems running a DTN protocol implementation designed and written at JPL for use on spacecraft, and further configured for higher bandwidth performance. Comparing baseline data transmission metrics with and without periodic optical link interruptions, the experiment confirmed the DTN protocols’ ability to handle real-world unexpected link outages while maintaining capability of reliably delivering data at relatively high rates. Finally, performance characterizations from this data suggest performance optimizations to configuration and protocols for future optical-specific DTN space link scenarios.

Keywords—Optical communications; Space networking; Disruption tolerant network; Interplanetary Overlay Network

I. INTRODUCTION

As an emerging and enabling communications technology, disruption-tolerant networks (DTNs) are capable of enabling reliable data exchange in challenging environments. As designed, DTNs can become a network of networks, bridging centers or outposts of rich connectivity across regions sparse

emptiness. The communications links required to transport data between separated networks can suffer degradation from numerous sources, and are often further constrained by power or infrastructure limitations at one or both ends. Optical communications become a natural fit, and a conceptual cornerstone of future space networking. Yet, signal fades in the optical link caused by clear air turbulence and clouds result in loss of data [1]. The exploration herein – an initial link prototype combining laser communications with DTN protocols – explores this operational concept and presents initial results on the behavior of such a system.

II. EXPERIMENT

A. System Description

Application of DTN protocols to a solar system internet and flight experiments on the Deep Impact spacecraft have been described in the open literature [2, 3]. These examinations have led to the investigation of DTN technology in the field of optical communications – where high-throughput links are subject to outages caused by transient obscuration such as cloud cover. Initial laboratory testing of the DTN protocols on an optical link was conducted with systems assembled in a point-to-point topology. The DTN protocol software used was JPL’s Interplanetary Overlay Network (ION), designed and written for use in deep space or similarly constrained computing environments. The physical optical link was created using two Perle media converter systems interfaced with a standard Ethernet switch. A schematic of the configuration is shown in Fig. 1 below. Relevant characteristics included a high rate

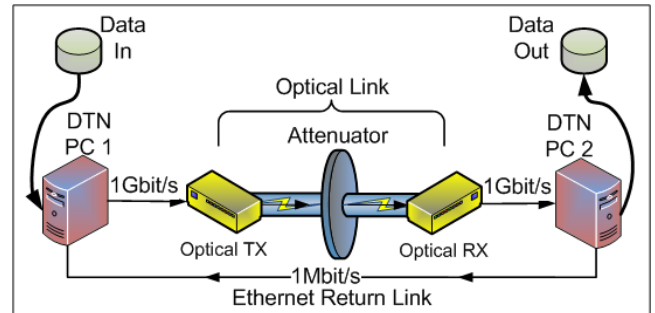


Figure 1. System schematic of DTN and optical hardware configuration.

forward optical link (fiber and free-space) and a low rate backchannel copper Ethernet link. A photo of the optical portion of the test configuration is provided in Fig. 2. The transmitter output was collimated and coupled to the focusing lens of the receiver fiber. The free space optical beam passed through a variable rotating attenuator (0 – 10dB attenuation) and was further attenuated by a 3dB neutral density filter shown in Fig. 2. The system was designed so that at maximum attenuation the signal level was firmly below that required to close the link, and while at minimum attenuation the signal was clean and strong.

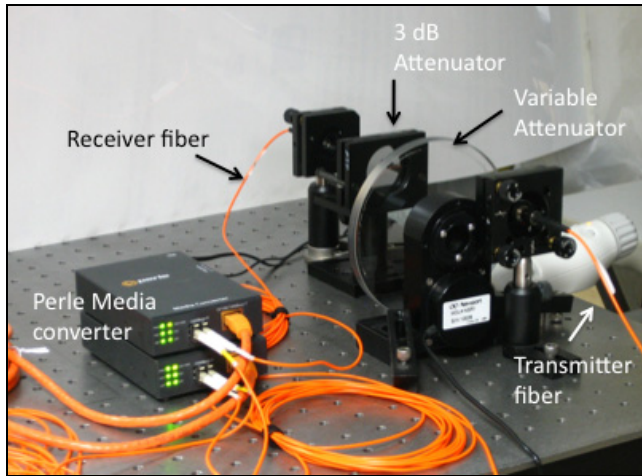


Figure 2. Free space optical link integrated with DTN systems.

The asymmetric data rate arrangement was chosen to mimic the operating concept of a unidirectional optical communications scenario, such as a downlink from a lunar or near-earth outpost. Fig. 3 illustrates the DTN protocol configuration, including the use of Licklider Transmission Protocol (LTP) for reliability below the bundle layer. The Datagram Retransmission (DGR) implementation of LTP in ION is a high-bandwidth, multithreaded solution that merges DTN protocols with an adaptive timeout interval congestion control method. The method is similar to that of internet TCP, but without the characteristic TCP ‘chattiness,’ allowing DGR to function in a way appropriate to the test scenario.

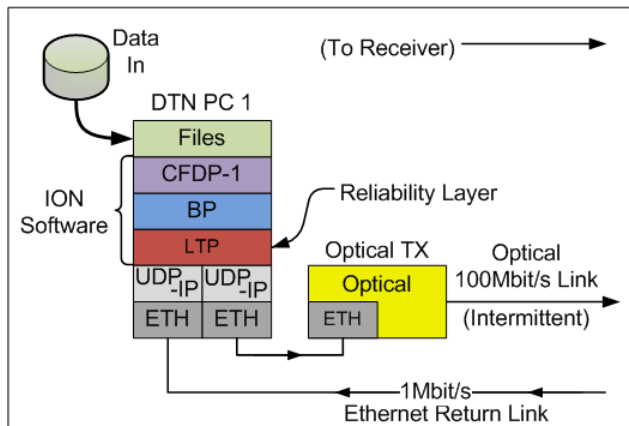


Figure 3. Transmitter DTN protocol diagram. The receiver protocol design (not shown) is identical in protocol layering.

Connections to and from the optical nodes and the backchannel link were achieved with 1Gbit/s Ethernet connections wired through a managed network switch. DTN PC 1 and 2 in Fig. 1 ran a Linux operating system and each used two separate Ethernet interfaces to establish the directional link architecture. A third network interface (not shown) was used for experiment control. Throughput on the backchannel link was hard-limited at the internet protocol (IP) network layer on both DTN PCs by modifying Linux kernel network queuing settings. A combination of IP routing and network layer 2 options configured in the Ethernet switch created segregated virtual local area networks (VLANs) and enforced the directionality of each link, as shown in the network diagram in Fig. 4.

Prior to testing, a network performance characterization

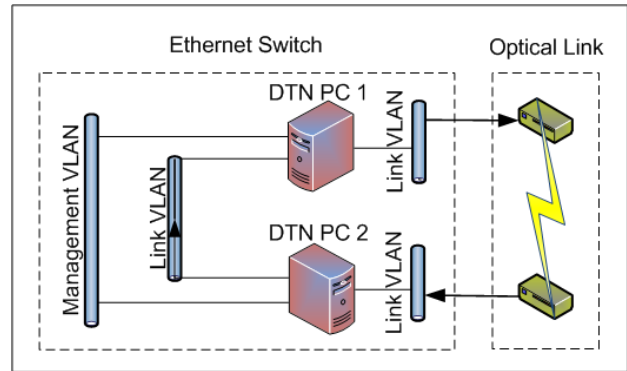


Figure 4. Ethernet network diagram.

software package was used to confirm basic network connectivity, performance and link directionality at the IP layer. Maximum TCP throughput was measured at approximately 800Mbit/s on the forward link, and 1Mbit/s on the return link, as expected.

B. Test Procedure

The investigation proceeded in two phases: (i) a fiber optic link between the two Perle media converters, and (ii) a free space optical link. To set a baseline for the overall system performance without link disruption, the fiber optic link was first compared to a copper Ethernet link, and a sweep of backchannel rates was performed. After this stage was complete, the optical signal was taken into free space and its performance was characterized with and without attenuation.

Each system test proceeded by generating data at DTN PC 1, packing it into bundles inserted into LTP segments and transmitting to a media converter via standard gigabit Ethernet. The data were then transmitted through the optical channel to the other media converter, and sent to DTN PC 2 again over standard gigabit Ethernet. All data was then extracted from segments and bundles and checked at DTN PC 2, which then sent acknowledgement of the received bundles through the low bandwidth Ethernet back channel to DTN PC 1. For missing data, retransmission requests were also sent. Throughout the tests, data was retained at PC 1 until PC 2 acknowledged receipt. For tests with sporadic link outages, the attenuator was operated at 1.7 rpm with random starting intervals for data transmission.

C. Data Analysis

The first set of tests compared optical forward and copper forward links, specifically looking at the performance of the forward link for a variety of return/back channel rates. This helped to establish parity between the two media and to characterize throughput dependence on backchannel size. Two subtypes of tests were used – (i) a burst of 30 million bytes to examine decreasing backchannel rates (Fig. 5), and then (ii) a steady-state (continuous) transfer with decreasing back channel rates (Fig. 6). The ‘bursty’ nature of the transfer of the 30 million bytes partially bypassed the congestion control mechanism of DGR, and as a result introduced some variation that averaged out in longer test intervals. Bundle size was fixed at 30,000 bytes in both cases. In terms of throughput and backchannel sensitivity, results from both media exhibited essentially similar performance.

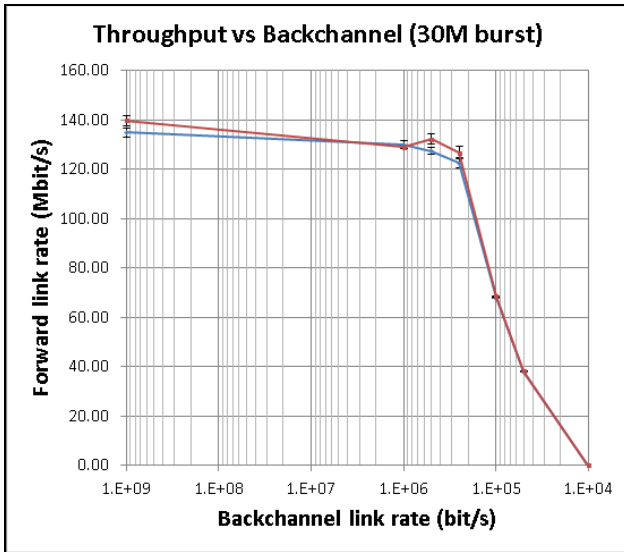


Figure 5. Copper (red) and fiber optic (blue) comparison for changing backchannel rates, bursting data.

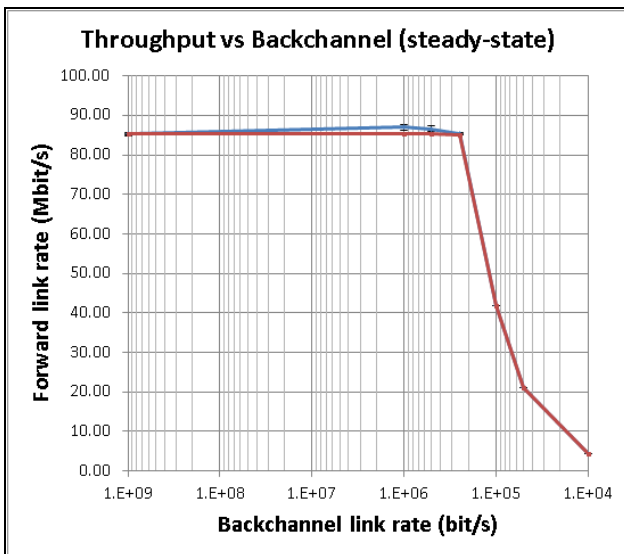


Figure 6. Copper (red) and optical (blue) comparison for changing backchannel rates, continuous data.

Next, the backchannel rate was fixed at 1Mbit/s and the bundle size varied. The optical signal was transmitted through free space, and a test series was run without link outages to confirm baseline, and again with the variable attenuator operational to disrupt the link. A total data aggregate of 30 million bytes (Fig. 7) and 60 million bytes (Fig. 8) was transmitted in bundles that varied in size (1,000 bytes to 60,000 bytes) and correspondingly in number. During the ‘on’ series of tests, the attenuator was not stopped or reset at any time, and ran continuously throughout the process.

This set of data was taken in a series of four tests for each bundle size. These tests were automatically sequenced and were not synchronized with the rotation of the attenuator, which remained fixed at 1.7 rpm. Although the attenuator loss was periodic and repeatable, the overlap with data transmission from test to test was uncorrelated – the time between tests was

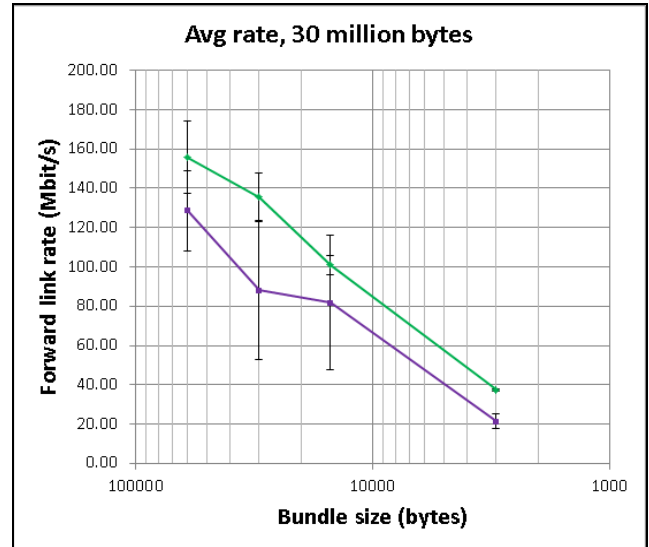


Figure 7. Attenuator on (green) and off (purple) comparison for 30 million bytes. Standard error was via four tests at each bundle size.

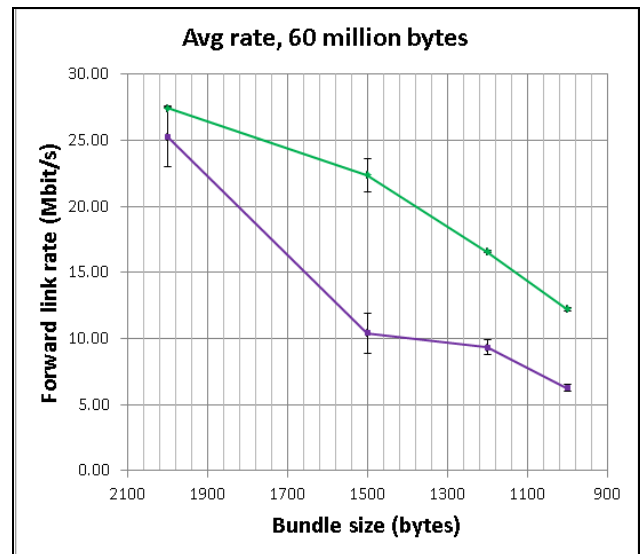


Figure 8. Attenuator on (green) and off (purple) comparison for 60 million bytes. Standard error was found between four tests at each bundle size.

random, and as such individual test duration was not deterministic. The end result was a random link loss pattern as would be expected in an optical link: Figs. 7 and 8 show that with the attenuator rotating, the variation in throughput data rate was much greater for the larger bundle sizes where link interruptions would have the potential for greatest disruption.

The 60-million byte tests in Fig. 8 were performed across a reduced range of small bundle sizes to show system performance when data was spread over more bundles and sent over a longer period of time. In this case, a random outage event affected a relatively small amount of data. The cost, however, was high: the increase in bundles also increased overhead in both data and processing time, especially given that this configuration utilized a 1:1 correspondence between DTN bundles, LTP blocks, LTP segments and UDP datagrams. It should be noted, however, that while this system operated in this fashion for simplicity in configuration and analysis, these ratios are tunable.

In general, the operation of the attenuator lowered the average data rate for every test of bundle size and count. This

behavior is shown in Figs. 9 and 10 with the time domain performance of a fixed bundle quantity and size (60,000 bundles, 1,000 bytes each, the attenuator on and off) producing visualizations of the data rate over the course of each individual test. To transmit the same data volume, it took approximately thirty-nine seconds with the attenuator off and sixty-six seconds with the attenuator enabled.

Several functional details in DGR are observable from these plots. Most clearly, the congestion control mechanism operates at almost regular intervals, severely chopping the throughput when a LTP segment is lost, and quickly rebuilding the data transfer rate back up if no further interruption occurs. Relatedly, true link outage periods required a longer recovery time for the system to ramp throughput back up after the reconnection of the link. These behaviors are artifacts of the TCP-like nature in the design of DGR's congestion control algorithms. In the first case (no link outages), the behavior aided the performance of the system, but in the second (attenuator on), it was a somewhat of a hindrance. From this data, it is clear that DTN protocols can properly handle

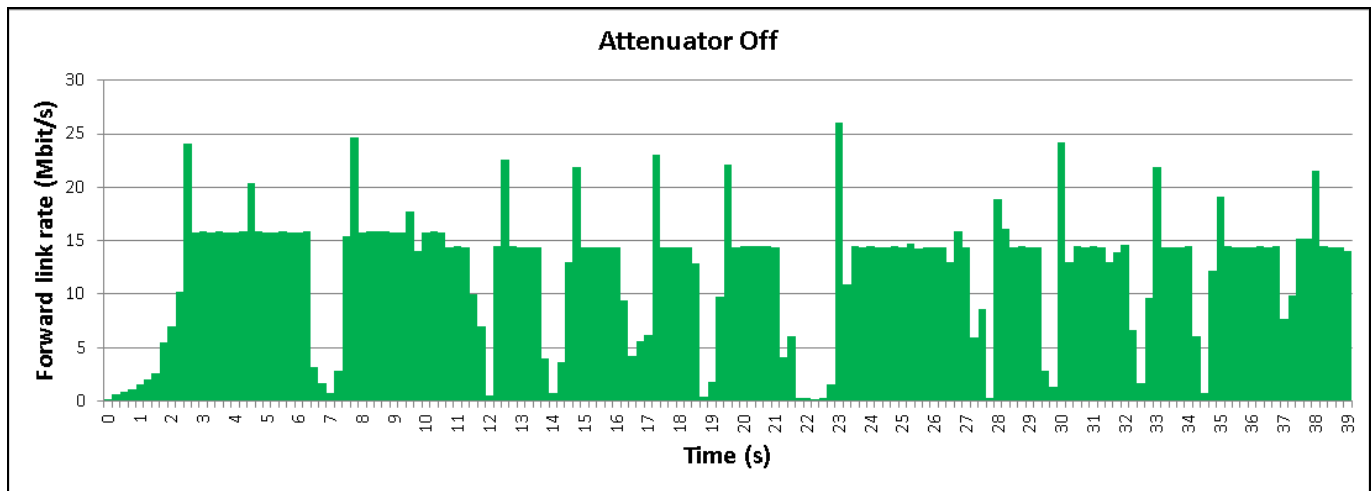


Figure 9. Attenuator off, 39 seconds of rate-time history for 60 million total bytes over 60,000 bundles.

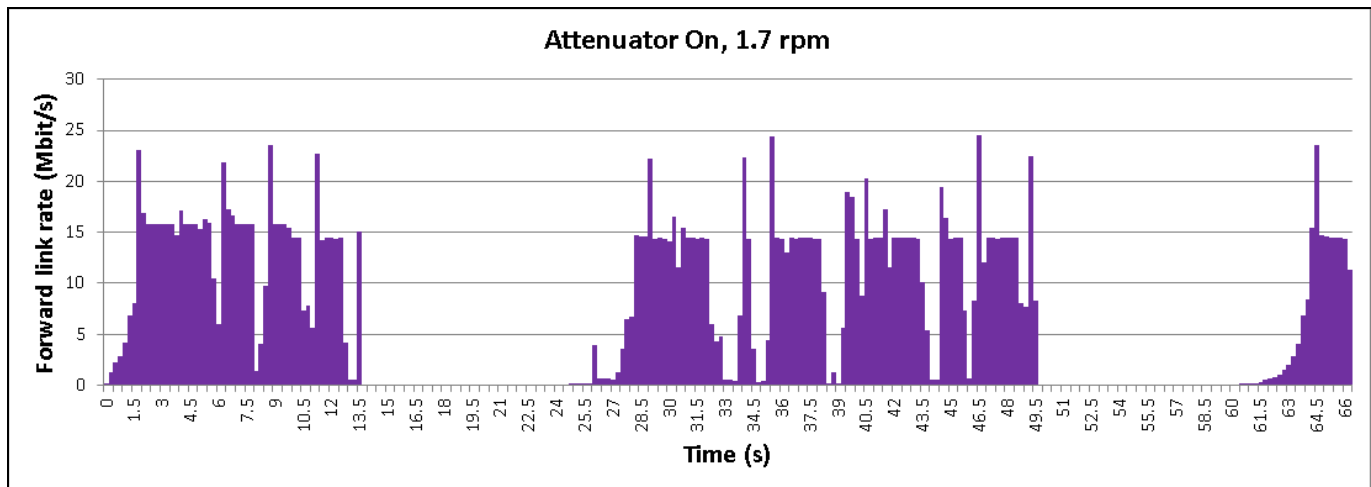


Figure 10. Attenuator on, 66 seconds of rate-time history for 60 million total bytes over 60,000 bundles.

unexpected outages, even at higher data rates. The LTP DGR system in particular has the necessary capability to maximize stable throughput in such a system, and is able to maintain data integrity even with unexpected outages when there is a low-latency (and, to a degree, low bandwidth) feedback channel.

There are a number of optimizations that could be made given a concerted effort toward further developing DTN protocols for this scenario. For example, a capacity to predict periodic link outages (or outage durations) could potentially improve recovery time and generally increase throughput. While the throughput of the network elements at the IP layer is quite high (near maximum for the hardware), processing delays in ION on the individual DTN PCs appeared to cause the occasional dropped UDP datagram, triggering congestion control throughput throttling. A hardware implementation of the DTN protocols could completely remove this performance hurdle and fully utilize all available bandwidth.

III. CONCLUSIONS

In this paper, we have reported results on the first experiments that demonstrate space DTN protocols and implementations over a free-space optical link. Initial findings show potential for excellent system performance, and also suggest that the architecture could be optimized to take advantage of the operational concepts inherent in optical communications – namely using link state information to burst high rates of data at re-acquisition instead of utilizing a rate ramping algorithm. Future work may include a novel approach to optimizing protocol behavior, coupled with a complete optical system (transmitter, receiver, and laser power) to ensure

that a link would support the required increase in data rate to transmit the buffered data corresponding to a representative outage period.

ACKNOWLEDGMENT

The work described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology under a contract with the National Aeronautics and Space Administration (NASA). Reference herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government, NASA or the Jet Propulsion Laboratory, California Institute of Technology.

REFERENCES

- [1] A. Biswas, K. E. Wilson, S. Piazzolla, J. Wu, W. H. Farr, "Deep-Space optical communication link availability and data volume," *Proceedings of SPIE, Free-Space Laser Communication Technologies XVI*, [Ed. G. S. Mecherle, C. Y. Young and J. S. Stryjewski], Volume 5338, 174, 2004.
- [2] E. Wyatt, E. Barkley, S. Burleigh et al, "Enabling Autonomous Exploration via the Solar System Internet," *IEEE Intelligent Systems*, vol. 25, no. 5, pp. 9-15, Sep./Oct. 2010.
- [3] J. Schoolcraft, S. Burleigh, R. Jones, E. Wyatt, J. Torgerson, "Deep Impact Network Experiments - Concept, Motivation and Results," *Proc. Space Operations 2010 Conf.*, AIAA, 2010.
- [4] S. Burleigh, "Interplanetary Overlay Network (ION) Design and Operation," JPL D-48259, V1.10, October 2009.